

Cross-Validated Wavelet Block Thresholding for Non-Gaussian Errors

K.McGinnity^a, R.Varbanov^{b,*}, E.Chicken^b

^a*Institute for Defense Analyses, Alexandria, Virginia*

^b*Department of Statistics, Florida State University, Tallahassee, FL*

Abstract

Wavelet thresholding generally assumes independent, identically distributed normal errors when estimating functions in a nonparametric regression setting. VisuShrink and SureShrink are just two of the many common thresholding methods based on this assumption. When the errors are not normally distributed, however, few methods have been proposed. A distribution-free method for thresholding wavelet coefficients in nonparametric regression is described, which unlike some other non-normal error thresholding methods, does not assume the form of the non-normal distribution is known. Improvements are made to an existing even-odd cross-validation method by employing block thresholding and level dependence. The efficiency of the proposed method on a variety of non-normal errors, including comparisons to existing wavelet threshold estimators, is shown on both simulated and real data.

Keywords: wavelets, thresholding, nonparametric function estimation

1. Introduction

Wavelet thresholding has been a staple of statistical functional estimation for years. Donoho and Johnstone (1994, 1995, 1998) introduced VisuShrink and SureShrink methods for thresholding the wavelet coefficients derived from the wavelet transformation of the observed data in nonparametric regression:

$$y_i = f(x_i) + \varepsilon_i, \quad i = 1, 2, \dots, n, \quad (1)$$

where the ε_i are independent and identically distributed (iid) Gaussian errors with mean 0 and constant variance σ^2 and the sample points $x_i = i/n$ are equally spaced over an interval. The assumptions on the errors have been loosened in only a handful of papers on wavelet thresholding.

Neumann and von Sachs (1995) discuss wavelet thresholding methods in non-Gaussian and non-iid situations. The main idea of their paper is that, in many situations, asymptotic normality can be proven and traditional thresholding methods can be used. Given independent observations, they demonstrate a way to show equivalence to the Gaussian case via strong approximations. They also derive asymptotic normality in the case of weak dependence.

Antoniadis and Fryzlewicz (2006) propose a scale-dependent wavelet thresholding procedure for Gaussian noise, and then extend it to include non-Gaussian noise. However, the paper assumes not only that the non-normal errors are iid with mean zero, but also that they follow a known specified distribution. They determine a suitable threshold for each resolution level by mimicking the arguments of Donoho and Johnstone in the Gaussian case.

Pensky and Sapatinas (2007) investigate the performance of Bayes factor estimators in wavelet regression models with iid non-Gaussian errors. They choose a general distribution η_j for the errors and assume they possess symmetric PDFs on the reals that are unimodal,

*Corresponding author

Email address: r.varbanov@stat.fsu.edu (R.Varbanov)

positive, and finite at zero. One advantage of their method is that knowledge of the true distribution of the errors is not needed in order to obtain an optimal estimator of f . However, their estimators are only preferable for irregular functions with high peaks, and produce sub-optimal results when compared with other methods under certain prior distributions.

Nason (1996) introduces an even-odd cross-validation method for choosing the threshold parameter in wavelet shrinkage. His statistic compares an interpolated wavelet estimator from the even reconstructed data to the odd noisy data and vice versa over various threshold values, then applies a sample size correction.

In this paper, we propose a completely nonparametric method to threshold wavelet coefficients that enhances Nason's cross-validation method by incorporating level-dependent block thresholding. Block thresholding thresholds wavelet coefficients in groups, rather than individually, with the goal of increasing precision by utilizing information about neighboring coefficients (Cai (1999)). Nason's method uses term-by-term thresholding, so it is reasonable to ask if incorporating blocking will have an analogous effect here.

Nason also makes use of a global threshold, the same threshold value for all considered coefficients. This is similar to VisuShrink of Donoho and Johnstone (1994). However, level-dependent thresholding has also been shown to have advantages over universal thresholds. For example, SureShrink (Donoho and Johnstone (1995)), a level-dependent thresholding method, has been shown to have lower MSE than VisuShrink. Each of these modifications, blocking and level dependence, improves performance with distribution-based thresholds and thus are natural considerations for attempting to improve cross-validation thresholding.

Our method does not put any assumptions on the errors except that they are iid and centered at zero. Unlike Neumann and von Sachs, we do not discuss asymptotic normality, but instead develop a method specifically meant to handle non-Gaussian errors. Nor do we require that the distribution of the errors be known, as do Antoniadis and Fryzlewicz. Unlike that of Pensky and Sapatinas, no proper choice of prior is required for our method.

This paper is divided as follows. Section 2 provides a brief background on wavelets, wavelet notation, and wavelet thresholding methods before the details of the proposed estimator are described in Section 3. Section 4 contains a simulation comparison of the proposed estimator to the Nason estimator, VisuShrink, and other current methods which assume normal errors, as well as two example applications of our method using inductance plethysmography and vertical profile density data. A discussion of the results and methods is given in the final section.

2. Background

2.1. Wavelets

Wavelets are an orthogonal series representation of functions in the space of square-integrable functions $L_2(\mathbb{R})$. Ogden (1997) and Vidakovic (1999) offer good introductions to wavelet methods and their properties. Let ϕ and ψ represent the father and mother wavelet functions, respectively. There are many choices for these two functions, see Daubechies (1992). Here, ϕ and ψ are chosen to be compactly supported and to generate an orthonormal basis. Let

$$\phi_{jk}(x) = 2^{j/2}\phi(2^jx - k)$$

and

$$\psi_{jk}(x) = 2^{j/2}\psi(2^jx - k)$$

be the translations and dilations of ϕ and ψ , respectively. For any fixed integer j_0 ,

$$\{\phi_{j_0k}, \psi_{jk} | j \geq j_0, k \text{ an integer}\}$$

is an orthonormal basis for $L_2(\mathbb{R})$. Let

$$\xi_{jk} = \langle f, \phi_{jk} \rangle$$

and

$$\theta_{jk} = \langle f, \psi_{jk} \rangle$$

be the usual inner product of a function $f \in L_2(\mathbb{R})$ with the wavelet basis functions. Then f can be expressed as an infinite series:

$$f(x) = \sum_k \xi_{j_0 k} \phi_{j_0 k}(x) + \sum_{j=j_0}^{\infty} \sum_k \theta_{jk} \psi_{jk}(x). \quad (2)$$

The function f is not known and must be estimated. This is done using the discrete wavelet transform (DWT) of Mallat (1999). If f is sampled as a vector of dyadic length $n = 2^J$ for some positive integer J , then the DWT will provide a total of n estimated coefficients $\xi_{j_0 k}$ and θ_{jk} over the indices $j = j_0, j_0 + 1, \dots, J - 1$ and for all appropriate k . The lowest level possible for j_0 is 0, the highest is $J - 1$.

The wavelet basis functions are easily periodized to a specified interval. In this paper, we use wavelets that have been periodized to the interval $[0, 1]$. In this case, the index k for resolution level j runs from 1 to 2^j in (2).

Wavelets have the useful property that they can simultaneously analyze a function in both time and frequency. This is done by projecting the function to be analyzed into several subspaces or resolution levels. Each resolution level represents a different degree of smoothness of the function. The lowest resolution level, associated with the index $j = j_0$, represents the smoothest or coarsest part of the function. Increasing the index j corresponds to decreasing smoothness. The highest resolution levels j therefore represent the behavior of the function at the highest frequencies or scales. Since the wavelet series (2) forms an orthogonal representation, the sum of the projections in these resolution levels is the original function f . The construction of wavelet functions ϕ and ψ provide the ability to localize the analysis within each subspace. The higher the resolution, the greater the degree of localization.

By varying the resolution level j , wavelets have the ability to zoom in or out onto the smooth or detailed structure of f . This is referred to as the multiresolution property of wavelets. Changing the index k allows wavelets to localize the analysis. These properties enable wavelets to model functions of very irregular types, as well as smooth functions.

We use W to denote the $n \times n$ DWT transformation matrix. Applying the DWT to the observed values $y = (y_1, y_2, \dots, y_n)'$ in (1) gives the estimated wavelet coefficients

$$\tilde{\theta} = \left(\tilde{\xi}_{j_0 1}, \tilde{\xi}_{j_0 2}, \dots, \tilde{\xi}_{j_0 2^{j_0}}, \tilde{\theta}_{j_0 1}, \tilde{\theta}_{j_0 2}, \dots, \tilde{\theta}_{J-1, 2^{J-1}} \right)' = W y.$$

Applying the inverse DWT $W^{-1} = W'$ to these coefficients returns the original data, $y = W' W y$.

2.2. Thresholding

Most wavelet analysis uses some form of thresholding. These can be term-by-term methods, where each individual wavelet coefficient is modified individually, or block methods, where coefficients are modified in groups. There are two general types of term-by-term thresholding methods. Hard thresholding is a “keep or kill” operation that sets to 0 any wavelet coefficient below a certain threshold value λ . In soft thresholding, coefficients smaller in magnitude than λ are set to 0 and coefficients larger than λ are shrunk towards 0.

A popular method of term-by-term thresholding is the VisuShrink method of Donoho and Johnstone (1994). For a single estimated coefficient $\tilde{\theta}_{jk}$,

$$\hat{\theta}_{jk} = \eta(\tilde{\theta}_{jk}, \lambda) = \text{sgn}(\tilde{\theta}_{jk})(|\tilde{\theta}_{jk}| - \lambda)_+$$

is the thresholded coefficient, where λ is a threshold parameter. This is an example of soft thresholding. Hard thresholding via Visushrink can also be used. The value of λ is chosen to

give optimal results in terms of reconstruction; $\lambda = \sigma\sqrt{2\log(n)}$ where σ and n are from (1). This is known as the universal threshold since the same threshold value is used across all resolution levels. A thresholded estimate is then formed as

$$\hat{f} = W'\eta(Wy, \lambda).$$

Note that only select resolution levels are subjected to thresholding. Typically these are the highest m resolution levels, where m is specified by the user.

Another popular thresholding method combines VisuShrink with a resolution-level-dependent method based on Stein's Unbiased Risk Estimator (SURE), see Stein (1981). This SURE based method obtains optimal mean squared error (MSE) but does not work well if the detail coefficients at a given resolution level are too sparse. Thus, this hybrid method will implement the universal threshold if a level is too sparse, and will use the SURE based method otherwise. Level j is too sparse if

$$2^{-j} \sum_{k=1}^{2^j} (d_{j,k}^2 / \sigma_{j,k}^2 - 1) \leq 2^{-j/2} (\log_2 2^j)^{3/2}.$$

One may choose to ignore the sparsity condition. We will refer to this as SureShrink in our paper, see Donoho and Johnstone (1995). Both the hybrid method (HybridShrink) and SureShrink are only designed to be used with soft thresholding, and are more computationally expensive than VisuShrink.

Block thresholding methods typically employ variants of the block projection estimator of Cai (1999). If B_i , $i = 1, \dots, N$ are blocks of equal size L ($\sim \log n$) that evenly divide the observed wavelet coefficients $\tilde{\theta}$, \mathcal{H} is a subset of the block indices $\{1, \dots, N\}$, and $\tilde{\theta}_{B_i}$ are the L coefficients in block B_i , the goal is to have \mathcal{H} consist only of those blocks where the signal is greater than the noise:

$$\mathcal{H} = \mathcal{H}(\tilde{\theta}) = \{i : \|\tilde{\theta}_{B_i}\|_2^2 > L\sigma^2\}, \quad (3)$$

where $\|\tilde{\theta}_{B_i}\|_2^2 = \sum_{k \in B_i} \tilde{\theta}_k^2$.

Anything less than the threshold is set to zero, giving the final block projection estimator to be

$$\hat{\theta}_{B_i} = \begin{cases} \tilde{\theta}_{B_i} & \text{if } i \in \mathcal{H} \\ 0 & \text{if } i \notin \mathcal{H}. \end{cases} \quad (4)$$

So, in any particular block, either the original coefficients are kept, or every coefficient in the block is set to zero.

In more recent advances, Barber and Nason (2004) build on the multiwavelet shrinkage method introduced by Downie and Silverman (1996) by implementing the use of a complex wavelet transform to estimate real signals. The Complex Multiwavelet Style (CMWS) shrinkage method uses the complex-valued Daubechies wavelets used by Lawton (1993) and Lina and Mayrand (1995) and discards the imaginary component of the reconstructed signal. For complex-valued empirical wavelet coefficient $d_{j,k}^*$, where $d_{j,k}^* \sim N_2(d_{j,k}, \Sigma_j)$, a “thresholding statistic”

$$\theta_{j,k} = d_{j,k}^{*T} \Sigma_j^{-1} d_{j,k}^* \quad (5)$$

is computed. Hard thresholding is then carried out according to the rule

$$\hat{d}_{j,k}^{MH} = d_{j,k}^* I(\theta_{j,k} > \lambda), \quad (6)$$

where I is the indicator function and $\lambda = 2\log n$, a threshold derived by Downie and Silverman (1996) for use in the multiwavelet case.

Fryzlewicz (2007) combines thresholding and the unbalanced Haar (UH) basis introduced by Girardi and Sweldens (1997), for which jumps in the basis functions do not necessarily occur in the middle of their support. Each basis vector for the discrete UH transform is selected to best match the data at a specific scale and location by applying the idea of a *matching pursuit* algorithm originally introduced by Mallat and Zhang (1993). The author focuses on a *top-down* version of algorithm where the greatest concentration of the power of signal is placed on the coarse scales before proceeding to finer scales. After transformation, hard thresholding can be carried out by comparing the absolute value of empirical wavelet coefficient $Y_{j,k}$ to the universal threshold used in VisuShrink, $\lambda = \sigma\sqrt{2\log n}$.

Johnstone and Silverman (2005) propose the selection of a level-dependent threshold utilizing Empirical Bayes methods. Acknowledging the need for a threshold to adapt to the sparsity of a signal and consequently the resolution level, this Empirical Bayes method models this sparsity by selecting a suitable prior distribution for the wavelet coefficients of the signal. At each resolution level, the weights of the distributions that comprise the mixture prior are estimated using marginal-maximum likelihood. A threshold is then calculated such that the posterior median of the distribution of the coefficients is zero if and only if the magnitude of the coefficients is less than said threshold. Hard thresholding can then be applied using the selected threshold level.

It is important to note that all the methods mentioned above put strong assumptions on the errors, particularly normality.

The properties and advantages of differing types of thresholding methods are well documented in the literature. See Donoho and Johnstone (1994, 1998); Cai (1999); Pensky (1999); Cai and Silverman (2001); Chicken (2003, 2005).

3. The Estimator

We observe a vector of data, $y_i = f(x_i) + \epsilon_i$, $i = 1, \dots, n$ for which the noise ϵ_i is independent and identically distributed. If the noise is not centered at zero, we subtract the mean of the noise from each ϵ_i to eliminate any potential bias. Specific choices for the noise will be discussed in more detail later. When the discrete wavelet transform (DWT) is applied to the data, we obtain $Wy = Wf(x) + W\epsilon$ where W represents the wavelet transformation. If the original errors ϵ have mean 0 and covariance $\Sigma = \sigma^2 I_n$ then,

$$E(W\epsilon) = 0$$

and

$$\text{Cov}(W\epsilon) = W\Sigma W' = W\sigma^2 I_n W' = \sigma^2 W I_n W' = \sigma^2 W W' = \sigma^2 I_n.$$

Since W is an orthogonal transformation, these new errors $W\epsilon = \epsilon'$ are uncorrelated, as shown above by the fact that the covariance matrix is diagonal. They are not necessarily independent unless the ϵ_i have a Gaussian distribution. Opsomer et al. (2001) show that issues in nonparametric regression only arise with *correlated* errors. The errors are also identically distributed within each resolution level (see McGinnity and Chicken (2012)).

Our approach is motivated by Nason's even-odd cross-validation procedure. We propose three major modifications. First, we make the threshold level-dependent, meaning we allow a different threshold value to be chosen for each resolution level, rather than a global threshold for all resolution levels. This provides increased estimability and accuracy. For example, SureShrink, which has a level-dependent threshold, has lower MSE than VisuShrink, which has a global threshold, but SureShrink often has noisier reconstructions. Second, we employ block thresholding. Block thresholding divides the data into blocks of neighboring coefficients, and will set to zero an entire block if the sum of the squared coefficients in that block is less than the threshold value. Block thresholding has been shown to have improved error and visual fit over traditional term-by-term

thresholding (Cai (1999)), alleviating the above concern about “visually pleasing” reconstructions vs. low MSE. Third, we add an additional term to the error statistic that gets minimized during cross-validation. This term compares two sampled reconstructions to one another.

To set up our algorithm, we first divide the noisy data of length n into even and odd parts, y^e and y^o :

$$y^e = (y_2, y_4, \dots, y_n)'$$

and

$$y^o = (y_1, y_3, \dots, y_{n-1})'.$$

We then apply the DWT to each part to obtain \hat{f}^e and \hat{f}^o . We choose the block size to be the closest dyadic number to $\log(n)$ (shown to possess optimal properties in Cai (1999)). We define the threshold search range to be zero to the maximum sum of squared coefficients in any block. This is the most conservative search range we can choose since the minimum threshold will set all detail coefficients to be zero, giving the smoothest possible reconstruction, and the maximum threshold will set nothing to zero and return the original noisy data. It is possible to reduce the size of this search range through methods such as pilot density estimation on coefficients in the first resolution level, however, in doing so, one risks leaving out the optimal threshold value from the search.

In the first iteration, we start by finding the optimal threshold value λ for all detail resolution levels simultaneously by minimizing a combination of two errors. The first error is the sum of the squared difference between the odd reconstruction and the even data, and the squared difference between the even reconstruction and the odd data. This mimics the method of Nason (1998), except we apply the threshold rule to all coefficients in a block rather than individual coefficients. The second error is the sum of the squared difference between the odd reconstruction and the even reconstruction. This error is not considered by Nason. We choose to use this combination of errors in order to obtain the optimal smoothness; intuitively, we expect a comparison of the noisy data to the smooth reconstruction to under-smooth our final estimate and a comparison solely between the even and odd reconstructions to over-smooth our final estimate. Thus, the initial global threshold estimate is

$$\hat{\lambda} = \underset{\lambda}{\operatorname{argmin}} \left\{ \frac{1}{2} \{ \|\hat{f}_\lambda^o - y^e\|_2^2 + \|\hat{f}_\lambda^e - y^o\|_2^2 \} + \frac{1}{2} \{ \|\hat{f}_\lambda^o - \hat{f}_\lambda^e\|_2^2 \} \right\}. \quad (7)$$

We then proceed to make our estimator level-dependent. Letting λ_j be the threshold value for resolution level j ($j = j_0, \dots, J-1$), we fix λ_{J-1} at the value found in (7) and repeat the search described above to find an optimal threshold for resolution levels $J-2, \dots, j_0$. This value becomes our λ_{J-2} and we next search for the optimal threshold for resolutions levels $J-3, \dots, j_0$. This value becomes our λ_{J-3} and we then find the optimal threshold for resolution level $J-4, \dots, j_0$. Repeat this process until the optimal λ_{j_0} for j_0 alone is found. These serve as our initial level-dependent threshold estimates.

In the second and all subsequent iterations, we modify each resolution level threshold value one at a time until convergence is achieved. First, fix $\lambda_{J-2}, \lambda_{J-3}, \dots, \lambda_{j_0}$ at their previous iteration values and search for the optimal λ_{J-1} . Next, fix λ_{J-1} at this new value and $\lambda_{J-3}, \dots, \lambda_{j_0}$ at their previous iteration values and search for the optimal λ_{J-2} . Then, fix λ_{J-1} and λ_{J-2} at their new values and $\lambda_{J-4}, \dots, \lambda_{j_0}$ at their previous iteration values and search for the optimal λ_{J-3} . Continue until all λ_j have been modified individually. We found that this process generally only had to be performed once before all threshold values converged.

The last step in our procedure is also a modification of Nason’s final step. Since all estimates in his paper were based on $n/2$ data points, a heuristic method based on the universal threshold was used to obtain a threshold for n data points. Once his error statistic had been minimized to obtain a threshold value $\lambda^{n/2}$, the following correction was applied to obtain the final threshold:

$$\lambda^n \approx \left(1 - \frac{\log 2}{\log n}\right)^{-1/2} \cdot \lambda^{n/2}. \quad (8)$$

where the superscript refers to the number of points used to obtain the threshold value.

In our case, the sample size used to calculate the statistic changes depending on which resolution level j we are working with ($j = 1$ representing the highest resolution level). Thus we apply the following level-dependent correction:

$$\lambda_j^{n/2^j} \approx \left(1 - \frac{\log 2}{\log(\frac{n}{2^j})}\right)^{-1} \cdot \lambda_j^{n/2^{j+1}}, \quad j = 1, 2, 3, 4, \dots, J - j_0 \quad (9)$$

Notice that we no longer need to take the square root in our correction since our threshold values are based on the sum of *squared* coefficients within a block. Again, the superscript of λ refers to the number points used to obtain the threshold value and does not indicate an exponent.

4. Simulation and Examples

4.1. Simulation

Our estimator is compared to the traditional thresholding method of VisuShrink, Nason's method, the CMWS shrinkage of Barber and Nason (2004), the UH technique of Fryzlewicz (2007), and the Empirical Bayes (EBayes) selection of thresholds of Johnstone and Silverman (2005) via mean squared error (MSE). (See McGinnity and Chicken (2012) report for comparisons to Antoniadis' method.)

Using sample sizes ranging from 2^9 to 2^{11} and signal to noise ratios (SNRs) of 3 and 5, these methods were tested on the eight standard test functions (Blip, Blocks, Bumps, Corner, Doppler, Heavisine, Spikes, and Wave), see Donoho and Johnstone (1994) and Marron et al. (1998). Figure 1 shows a plot of these functions.

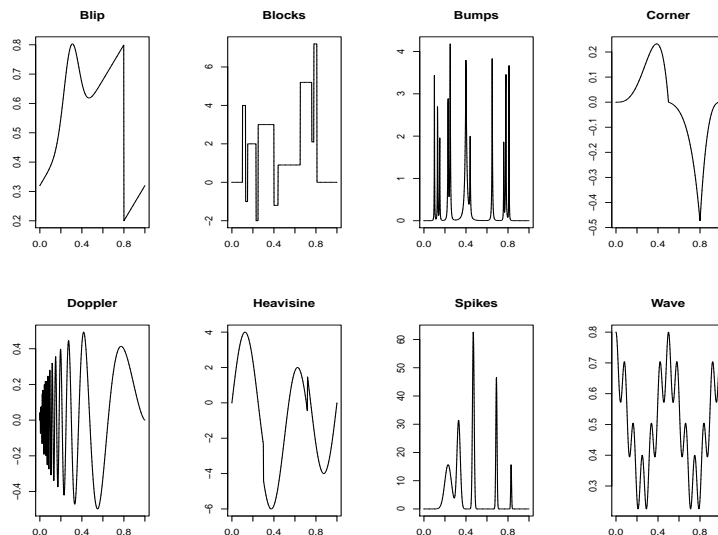


Figure 1: The eight test functions used in this paper sampled at 512 points.

Various types of noise were added on top of these test functions to create our data. The only assumptions imposed on the errors were that they were independent and identically distributed. In this paper we show results for T_3 and lognormal errors. Other types of noise were considered, but these two were chosen to exemplify a heavy-tailed distribution and a skewed distribution.

The noise was scaled appropriately to obtain the desired SNR, thus the true distribution of the noise was not necessarily known. For example, a scaled t -distribution is no longer a t . In the case of skewed distributions such as lognormal, we first centered the noise at the mean before applying it to the test function. Figures 2 and 3 show plots of T_3 and lognormal noise added to all the test functions.




Figure 2: The test functions with scaled T_3 noise added to represent a SNR of 5.




Figure 3: The test functions with centered lognormal noise added to represent a SNR of 5.

When possible, all functions were analyzed using the least asymmetric wavelet basis of length eight (LA-8) (Daubechies (1992)) with periodic boundary handling, and the lowest resolution level, j_0 , set to $J - 4$. For the CMWS method, estimation was done with all available complex-valued wavelets and the results were averaged. For the UH method, the functions were analyzed with the best top-down UH decomposition. For the EBayes method, a Laplace prior was used for the wavelet coefficients at each level. All simulations were performed in R (R Development Core

Team (2009)) using the *waveslim* (Whitcher (2009)), *WaveThresh* (Nason (2013)), *unbalhaar* (Fryzlewicz (2010)), and *EBayesThresh* (Silverman (2010)) packages.

Tables 1, 2, 3 and 4 show simulation results of 100 repetitions each for T_3 and lognormal noise respectively. Hard thresholding and various combinations of n , SNR, and test function were used. The numbers in the table represent the ratio of the average MSE for the given method to the average MSE of VisuShrink. Thus a ratio less than one indicates that a particular method outperformed VisuShrink, while a ratio greater than 1 indicates that VisuShrink performed better. In parentheses, the standard deviation ($\times 10^{-3}$) of the MSE from the 100 repetitions of the corresponding method is given. For each combination of conditions, we have identified the best performing method by highlighting the lowest average MSE. We also highlighted any other average MSE that was not significantly different than the lowest average MSE according to a paired t -test at the $\alpha = 0.05$ significance level.

Table 1: T_3 noise with SNR=5: Ratio of average MSE of listed method to average MSE of VisuShrink with corresponding standard deviation ($\times 10^{-3}$) of MSE of listed method

Function	n	LD Block	Nason	CMWS	UH	EBayes
Blip	512	0.60(0.02)	0.47(0.05)	0.94(0.21)	1.22(0.22)	1.35(0.23)
	1024	0.43(0.02)	0.49(0.04)	0.90(0.14)	1.14(0.15)	1.40(0.16)
	2048	0.29(0.02)	0.36(0.08)	0.91(0.14)	1.10(0.15)	1.44(0.15)
Blocks	512	1.45(19.32)	1.10(15.92)	0.79(24.05)	0.61(25.04)	1.08(25.95)
	1024	0.88(2.99)	0.85(5.92)	0.82(16.44)	0.71(17.11)	1.21(18.52)
	2048	0.57(2.42)	0.64(9.15)	0.82(16.19)	0.78(16.71)	1.28(17.36)
Bumps	512	1.68(2.04)	2.09(3.70)	0.65(2.37)	0.92(2.45)	0.99(2.54)
	1024	1.33(1.45)	2.12(3.47)	0.76(1.65)	1.18(1.76)	1.23(1.82)
	2048	0.84(0.72)	1.18(1.90)	0.83(1.64)	1.27(1.67)	1.33(1.73)
Corner	512	0.20(0.01)	0.25(0.03)	0.93(0.14)	1.35(0.14)	1.34(0.16)
	1024	0.19(0.01)	0.24(0.03)	0.92(0.09)	1.29(0.10)	1.41(0.11)
	2048	0.19(0.01)	0.25(0.05)	0.92(0.10)	1.18(0.10)	1.44(0.10)
Doppler	512	0.91(0.24)	0.83(0.22)	0.74(0.46)	1.50(0.49)	1.19(0.49)
	1024	0.42(0.11)	0.73(0.14)	0.83(0.31)	1.53(0.32)	1.34(0.35)
	2048	0.43(0.03)	0.47(0.17)	0.87(0.31)	1.41(0.31)	1.41(0.33)
Heavisine	512	0.21(4.28)	0.26(9.32)	0.94(48.41)	1.44(50.53)	1.35(53.36)
	1024	0.22(2.82)	0.26(8.60)	0.92(32.64)	1.33(34.44)	1.40(38.01)
	2048	0.20(2.77)	0.26(18.66)	0.92(32.66)	1.21(33.48)	1.44(35.04)
Spikes	512	1.09(345.62)	0.98(514.51)	0.81(555.43)	1.20(571.63)	1.24(601.57)
	1024	0.46(111.53)	0.56(145.44)	0.88(371.55)	1.31(386.95)	1.39(427.43)
	2048	0.22(37.16)	0.27(206.09)	0.90(370.63)	1.25(373.10)	1.43(398.18)
Wave	512	0.46(0.03)	0.47(0.02)	0.75(0.14)	1.39(0.15)	1.07(0.15)
	1024	0.21(0.01)	0.25(0.03)	0.92(0.09)	1.63(0.10)	1.41(0.11)
	2048	0.19(0.01)	0.25(0.05)	0.92(0.09)	1.41(0.09)	1.44(0.10)

Table 2: T_3 noise with SNR=3: Ratio of average MSE of listed method to average MSE of VisuShrink with corresponding standard deviation ($\times 10^{-3}$) of MSE of listed method

Function	n	LD Block	Nason	CMWS	UH	EBayes
Blip	512	0.43(0.03)	0.42(0.08)	0.95(0.35)	1.17(0.37)	1.34(0.39)
	1024	0.34(0.03)	0.39(0.06)	0.88(0.24)	1.09(0.25)	1.38(0.27)
	2048	0.26(0.02)	0.33(0.14)	0.91(0.24)	1.08(0.24)	1.43(0.26)
Blocks	512	0.98(18.83)	0.88(19.81)	0.81(39.75)	0.66(41.32)	1.12(42.93)
	1024	0.62(4.28)	0.63(7.50)	0.84(27.36)	0.76(28.71)	1.24(31.25)
	2048	0.44(2.55)	0.49(15.57)	0.85(27.04)	0.82(27.80)	1.32(29.10)
Bumps	512	1.31(3.25)	1.58(4.55)	0.68(3.92)	0.94(4.08)	1.01(4.25)
	1024	1.05(1.99)	1.61(4.26)	0.77(2.74)	1.13(2.91)	1.21(3.08)
	2048	0.65(0.73)	0.95(2.58)	0.84(2.73)	1.20(2.76)	1.33(2.89)
Corner	512	0.19(0.02)	0.24(0.05)	0.93(0.24)	1.30(0.24)	1.34(0.26)
	1024	0.19(0.01)	0.24(0.04)	0.92(0.16)	1.25(0.17)	1.41(0.18)
	2048	0.19(0.01)	0.25(0.09)	0.92(0.16)	1.16(0.16)	1.44(0.17)
Doppler	512	0.74(0.23)	0.70(0.24)	0.78(0.77)	1.41(0.80)	1.23(0.82)
	1024	0.38(0.14)	0.59(0.15)	0.84(0.52)	1.43(0.54)	1.34(0.60)
	2048	0.34(0.05)	0.39(0.30)	0.88(0.52)	1.35(0.52)	1.42(0.55)
Heavisine	512	0.20(7.14)	0.25(15.53)	0.93(80.76)	1.37(84.62)	1.35(89.06)
	1024	0.21(4.70)	0.25(14.34)	0.92(54.36)	1.29(56.92)	1.40(63.30)
	2048	0.20(4.62)	0.26(31.10)	0.92(54.45)	1.18(55.46)	1.44(58.46)
Spikes	512	0.86(377.77)	0.92(634.93)	0.83(923.13)	1.17(957.29)	1.26(1002.18)
	1024	0.43(142.21)	0.48(184.45)	0.89(618.77)	1.27(655.00)	1.39(714.08)
	2048	0.21(57.57)	0.27(352.48)	0.91(617.48)	1.22(619.83)	1.44(664.51)
Wave	512	0.36(0.04)	0.38(0.04)	0.81(0.23)	1.40(0.24)	1.18(0.25)
	1024	0.21(0.02)	0.25(0.04)	0.92(0.15)	1.55(0.17)	1.41(0.18)
	2048	0.19(0.01)	0.25(0.09)	0.92(0.15)	1.33(0.16)	1.44(0.17)

Table 3: Lognormal noise with SNR=5: Ratio of average MSE of listed method to average MSE of VisuShrink with corresponding standard deviation ($\times 10^{-3}$) of MSE of listed method

Function	n	LD Block	Nason	CMWS	UH	EBayes
Blip	512	0.37(0.04)	0.31(0.06)	0.97(0.20)	1.24(0.20)	1.34(0.21)
	1024	0.23(0.03)	0.31(0.07)	0.96(0.19)	1.21(0.19)	1.35(0.19)
	2048	0.17(0.02)	0.21(0.05)	0.96(0.13)	1.20(0.13)	1.39(0.14)
Blocks	512	1.23(18.94)	0.99(15.49)	0.90(22.09)	0.97(22.70)	1.18(23.23)
	1024	0.58(3.23)	0.62(8.11)	0.92(21.55)	1.03(21.72)	1.25(22.02)
	2048	0.39(2.36)	0.44(4.91)	0.93(14.80)	1.09(15.28)	1.33(15.89)
Bumps	512	1.56(2.28)	1.97(3.47)	0.83(2.15)	1.04(2.24)	1.13(2.27)
	1024	0.93(1.57)	1.55(3.96)	0.90(2.18)	1.15(2.17)	1.25(2.24)
	2048	0.58(0.72)	0.85(1.79)	0.93(1.49)	1.21(1.51)	1.34(1.59)
Corner	512	0.12(0.02)	0.17(0.03)	0.97(0.13)	1.27(0.13)	1.33(0.14)
	1024	0.11(0.01)	0.15(0.05)	0.97(0.13)	1.24(0.13)	1.35(0.13)
	2048	0.11(0.01)	0.14(0.03)	0.96(0.09)	1.22(0.09)	1.39(0.09)
Doppler	512	0.71(0.23)	0.67(0.22)	0.90(0.42)	1.32(0.43)	1.26(0.45)
	1024	0.25(0.11)	0.49(0.20)	0.94(0.41)	1.28(0.41)	1.31(0.42)
	2048	0.26(0.04)	0.29(0.09)	0.94(0.28)	1.27(0.29)	1.37(0.30)
Heavisine	512	0.13(6.26)	0.18(11.97)	0.97(44.70)	1.28(45.81)	1.32(47.74)
	1024	0.12(4.34)	0.17(16.17)	0.96(43.18)	1.24(43.51)	1.34(44.37)
	2048	0.11(2.87)	0.15(9.76)	0.96(29.57)	1.23(30.44)	1.39(32.14)
Spikes	512	0.75(331.14)	0.72(426.81)	0.92(508.73)	1.17(516.66)	1.25(532.81)
	1024	0.25(110.09)	0.36(197.79)	0.95(489.52)	1.21(492.82)	1.32(502.77)
	2048	0.13(34.40)	0.16(110.58)	0.95(336.12)	1.23(342.25)	1.38(363.60)
Wave	512	0.31(0.02)	0.35(0.03)	0.89(0.13)	1.24(0.13)	1.19(0.14)
	1024	0.11(0.01)	0.16(0.05)	0.97(0.12)	1.33(0.12)	1.35(0.13)
	2048	0.11(0.01)	0.14(0.03)	0.96(0.08)	1.29(0.09)	1.39(0.09)

Table 4: Lognormal noise with SNR=3: Ratio of average MSE of listed method to average MSE of VisuShrink with corresponding standard deviation ($\times 10^{-3}$) of MSE of listed method

Function	n	LD Block	Nason	CMWS	UH	EBayes
Blip	512	0.28(0.07)	0.29(0.09)	0.97(0.33)	1.23(0.33)	1.34(0.35)
	1024	0.20(0.04)	0.25(0.12)	0.96(0.31)	1.19(0.31)	1.34(0.32)
	2048	0.15(0.03)	0.19(0.07)	0.95(0.21)	1.19(0.22)	1.39(0.23)
Blocks	512	0.78(19.29)	0.76(19.78)	0.90(36.89)	0.98(37.80)	1.19(38.85)
	1024	0.39(3.92)	0.44(13.64)	0.92(35.87)	1.04(36.24)	1.25(36.77)
	2048	0.28(2.98)	0.32(8.18)	0.93(24.69)	1.09(25.39)	1.32(26.57)
Bumps	512	1.17(3.17)	1.44(4.22)	0.83(3.59)	1.03(3.75)	1.13(3.79)
	1024	0.76(2.64)	1.21(4.72)	0.91(3.64)	1.14(3.63)	1.25(3.72)
	2048	0.43(0.76)	0.67(2.11)	0.93(2.47)	1.20(2.52)	1.34(2.65)
Corner	512	0.12(0.03)	0.16(0.06)	0.97(0.22)	1.26(0.22)	1.33(0.23)
	1024	0.11(0.02)	0.15(0.08)	0.97(0.21)	1.23(0.21)	1.35(0.22)
	2048	0.11(0.01)	0.14(0.05)	0.96(0.14)	1.21(0.15)	1.39(0.16)
Doppler	512	0.54(0.20)	0.53(0.26)	0.90(0.70)	1.27(0.73)	1.25(0.75)
	1024	0.22(0.16)	0.40(0.27)	0.94(0.68)	1.26(0.69)	1.32(0.70)
	2048	0.20(0.06)	0.23(0.15)	0.94(0.47)	1.26(0.48)	1.37(0.51)
Heavisine	512	0.12(9.90)	0.17(19.91)	0.97(74.50)	1.28(76.21)	1.33(79.83)
	1024	0.11(7.23)	0.16(26.95)	0.97(72.00)	1.24(72.32)	1.35(74.13)
	2048	0.11(4.79)	0.15(16.26)	0.96(49.30)	1.22(50.74)	1.39(53.60)
Spikes	512	0.60(439.61)	0.67(529.10)	0.93(846.93)	1.17(870.70)	1.26(889.24)
	1024	0.24(152.73)	0.30(313.31)	0.96(815.74)	1.21(822.55)	1.33(839.48)
	2048	0.12(61.61)	0.15(184.31)	0.95(560.22)	1.22(572.79)	1.38(605.96)
Wave	512	0.23(0.03)	0.28(0.06)	0.91(0.21)	1.24(0.22)	1.22(0.23)
	1024	0.11(0.02)	0.16(0.08)	0.97(0.20)	1.30(0.21)	1.35(0.21)
	2048	0.11(0.01)	0.14(0.05)	0.96(0.14)	1.27(0.14)	1.39(0.15)

We considered three different classifications for the results of the proposed method (LD Block). Wins are cases where LD Block has the lowest average MSE and this average MSE is significantly different than all other competitors in pairwise comparisons. Losses are cases where LD Block has a significantly different average MSE than the leader. Ties are cases where LD Block is not significantly different than the leader or LD Block is the leader but is not significantly different than all other methods.

For T_3 noise and SNR=5, LD Block has 13 wins, 7 losses, and 3 ties. When SNR=3, LD Block has 16 wins, 4 losses, and 4 ties. Regardless of SNR, the proposed method loses only two of all the comparisons at the highest two sample sizes ($n = 2^{10}$ and 2^{11}). The only cases in which the proposed method is significantly worse than the leading method is for smaller sample sizes and test functions Blocks, Bumps, Doppler, and Spikes.

For lognormal noise and SNR=5, the proposed method has 18 wins, 3 losses, and 3 ties. The only time a competitor beats the proposed method in terms of MSE is at the smallest sample size and test functions Blips, Blocks, and Bumps. When SNR=3, LD Block has 20 wins, 1 loss, and 3 ties.

LD Block, Nason, and CMWS show substantial improvement over VisuShrink, UH, and EBayes for both T_3 and lognormal noise, particularly for smoother functions. This improvement increases as SNR decreases, as shown in the tables. Clearly, outside of CMWS, traditional methods that assume normality fail drastically when given heavy-tailed or skewed errors.

The proposed method also provides very visually appealing reconstructions for both smooth

and spiky functions. Figure 4 shows the average reconstruction over 10 repetitions for the proposed method and Nason’s method using the Blip test function, T_3 noise, SNR=3, and $n=512$. Both Nason’s and the proposed method provide much smoother average reconstructions than does VisuShrink (see McGinnity and Chicken (2012) for VisuShrink reconstruction), however, it is clear that, in four or five places, Nason’s method does not threshold a piece of data that it should have. This does not appear to happen at all for our level-dependent block method.




Figure 4: Average reconstructions for our method and Nason’s method using the Blip test function with T_3 noise, SNR=3, and $n=512$

Figure 5 shows the average reconstruction over 10 repetitions for the proposed method and Nason’s method using the Bumps test function, lognormal noise, SNR=3, and $n=1024$. While VisuShrink does not smooth the function enough and the reconstruction has a lot of excess noise (see McGinnity and Chicken (2012)), Nason’s reconstruction is overly smooth; it does not do a good job of picking up the jumps in the function and, in some cases, does not model them at all. The proposed method provides a “happy medium” visually. It models the spikes better than Nason’s method, but without the excessive noise seen in VisuShrink’s reconstruction. As a result, the MSE for our level-dependent block method is also the smallest of the three methods.




Figure 5: Average reconstructions for our method and Nason's method using the Bumps test function with lognormal noise, SNR=3, and $n=512$

We tested the robustness of our method by comparing the performance of our estimator to current estimators with Gaussian errors. Tables 5 and 6 show these results. Our method is intended for use in situations when errors are thought to be non-normal, therefore we do not necessarily expect our method to be an improvement. For normal noise and SNR=5, LD Block has 24 losses. For normal noise and SNR=3, LD block has 1 tie and 23 losses. With Gaussian errors, CMWS appears to be the optimal method as it has the lowest average MSE in all but 7 of the 48 cases. CMWS outperforms LD Block in all 48 cases with normal noise and has an average MSE that is on average 1.84 times lower than that of LD Block. However, in the 96 cases with non-normal errors, LD Block outperforms CMWS 81 times and has an average MSE that is on average 4.36 times lower than that of CMWS. We see that we outperform CMWS when we are expected to and we do so by a much wider margin than when CMWS outperforms our method.

Table 5: Normal noise with SNR=5: Ratio of MSE of listed method to MSE of VisuShrink with corresponding standard deviation ($\times 10^{-3}$) of MSE of listed method

Function	n	LD Block	Nason	CMWS	UH	EBayes
Blip	512	2.04(0.01)	1.07(0.02)	0.89(0.02)	1.38(0.02)	1.18(0.03)
	1024	1.20(0.02)	1.05(0.02)	0.68(0.01)	0.90(0.01)	1.00(0.02)
	2048	1.14(0.01)	1.00(0.01)	0.81(0.01)	0.81(0.01)	1.02(0.01)
Blocks	512	1.81(18.87)	1.05(6.06)	0.64(2.71)	0.13(2.91)	0.84(3.93)
	1024	1.40(1.88)	1.05(3.46)	0.68(1.82)	0.10(1.58)	0.87(2.36)
	2048	1.11(1.59)	1.04(1.72)	0.65(0.97)	0.08(0.98)	0.88(1.44)
Bumps	512	1.76(2.18)	1.94(2.08)	0.44(0.31)	0.86(0.69)	0.71(0.50)
	1024	1.72(1.29)	2.14(1.21)	0.50(0.20)	1.22(0.36)	0.84(0.34)
	2048	1.53(0.59)	0.94(0.22)	0.56(0.09)	1.50(0.17)	0.87(0.16)
Corner	512	0.98(0.01)	1.08(0.01)	0.86(0.01)	2.67(0.02)	1.29(0.02)
	1024	0.96(0.01)	1.07(0.01)	0.89(0.01)	2.03(0.01)	1.13(0.01)
	2048	0.97(0.00)	1.03(0.01)	0.87(0.00)	1.39(0.01)	1.06(0.01)
Doppler	512	1.54(0.19)	0.94(0.06)	0.46(0.04)	2.01(0.10)	0.95(0.09)
	1024	0.83(0.06)	0.95(0.04)	0.57(0.02)	2.38(0.05)	0.96(0.05)
	2048	1.39(0.02)	0.99(0.02)	0.67(0.02)	2.42(0.03)	0.98(0.02)
Heavisine	512	0.94(3.48)	1.06(4.34)	0.90(3.13)	2.89(5.60)	1.25(7.73)
	1024	0.96(1.99)	1.06(2.86)	0.87(1.97)	2.29(3.35)	1.12(3.40)
	2048	0.97(1.46)	1.03(1.78)	0.87(1.37)	1.61(2.12)	1.06(2.01)
Spikes	512	2.04(284.61)	0.95(62.18)	0.52(46.82)	1.40(82.13)	0.98(107.79)
	1024	1.41(87.44)	1.01(38.48)	0.68(23.64)	1.94(42.91)	1.04(45.91)
	2048	0.98(17.52)	1.03(20.41)	0.79(16.05)	1.97(25.73)	1.06(21.71)
Wave	512	0.97(0.01)	1.01(0.01)	0.45(0.01)	2.18(0.02)	0.82(0.02)
	1024	1.11(0.01)	1.06(0.01)	0.87(0.01)	4.56(0.01)	1.12(0.01)
	2048	0.97(0.00)	1.03(0.01)	0.87(0.00)	2.99(0.01)	1.06(0.01)

Table 6: Normal noise with SNR=3: Ratio of MSE of listed method to MSE of VisuShrink with corresponding standard deviation ($\times 10^{-3}$) of MSE of listed method

Function	n	LD Block	Nason	CMWS	UH	EBayes
Blip	512	1.47(0.02)	1.00(0.03)	0.87(0.03)	1.09(0.03)	1.13(0.06)
	1024	1.05(0.02)	1.01(0.02)	0.68(0.02)	0.75(0.02)	0.99(0.03)
	2048	1.11(0.02)	1.01(0.01)	0.84(0.01)	0.71(0.01)	1.03(0.01)
Blocks	512	1.44(18.82)	0.96(5.86)	0.67(4.00)	0.22(5.56)	0.90(6.58)
	1024	1.15(2.22)	0.99(3.47)	0.75(2.34)	0.17(2.71)	0.91(3.63)
	2048	1.02(1.50)	0.99(1.60)	0.73(1.42)	0.12(1.52)	0.96(1.94)
Bumps	512	1.54(2.79)	1.48(2.30)	0.51(0.54)	0.96(1.05)	0.79(0.78)
	1024	1.43(1.62)	1.43(1.55)	0.52(0.34)	1.09(0.57)	0.81(0.52)
	2048	1.28(0.55)	0.94(0.31)	0.58(0.15)	1.32(0.26)	0.86(0.26)
Corner	512	0.98(0.02)	1.08(0.02)	0.88(0.02)	2.34(0.03)	1.30(0.04)
	1024	0.96(0.01)	1.07(0.01)	0.89(0.01)	1.69(0.01)	1.13(0.02)
	2048	0.97(0.01)	1.03(0.01)	0.87(0.01)	1.18(0.01)	1.06(0.01)
Doppler	512	1.50(0.14)	0.99(0.08)	0.55(0.07)	2.03(0.14)	1.05(0.14)
	1024	0.83(0.09)	0.97(0.06)	0.61(0.04)	2.09(0.08)	0.97(0.07)
	2048	1.12(0.03)	1.00(0.03)	0.70(0.02)	2.08(0.04)	0.97(0.03)
Heavisine	512	0.93(5.31)	1.07(7.25)	0.90(5.20)	2.48(9.00)	1.27(13.05)
	1024	0.96(3.32)	1.06(4.75)	0.88(3.29)	1.95(4.91)	1.13(5.83)
	2048	0.97(2.43)	1.03(2.96)	0.87(2.28)	1.43(3.35)	1.06(3.36)
Spikes	512	1.89(365.21)	0.97(144.77)	0.55(71.82)	1.31(132.43)	0.99(165.34)
	1024	1.32(95.74)	0.99(61.06)	0.68(39.27)	1.67(67.18)	0.99(76.18)
	2048	0.98(30.43)	1.03(34.12)	0.83(26.45)	1.74(37.28)	1.06(35.89)
Wave	512	0.95(0.01)	1.03(0.02)	0.54(0.02)	2.29(0.03)	1.10(0.03)
	1024	1.08(0.02)	1.06(0.01)	0.87(0.01)	3.89(0.02)	1.13(0.02)
	2048	0.97(0.01)	1.03(0.01)	0.87(0.01)	2.60(0.01)	1.06(0.01)

We want to show that our comparisons are valid despite the selection of a different wavelet basis for two of the competitors (CMWS and UH). CMWS requires a complex-valued wavelet basis which cannot be applied to any of the other methods. However, the UH method selects an optimal Haar-like basis as its filter, so we used a Haar basis across the other methods to standardize comparisons. The results of these comparisons are shown for normal and non-normal noise in Tables 7 and 8, respectively. In the presence of normal noise, the standardization of the wavelet basis does not make UH the best thresholding technique as it only won or tied 7 of the 24 cases. This is fewer than the number of wins or ties by LD Block and EBayes, which had 8 and 9 such cases, respectively. In the presence of non-normal noise, which UH is not meant for, the gap between LD Block and UH is even greater as LD Block wins 18 of 24 cases while UH wins only 3. We acknowledge the fact that UH outperforms our method by a large factor when estimating the Blocks function in the presence of normal noise and attribute this to the fact that UH is designed to use a Haar basis, which is suited for estimating piecewise linear functions such as Blocks.

Table 7: Haar basis with normal noise and SNR=5: Ratio of MSE of listed method to MSE of VisuShrink with corresponding standard deviation ($\times 10^{-3}$) of MSE of listed method*

Function	n	LD Block	Nason	UH	EBayes
Blip	512	2.58(0.10)	1.39(0.01)	1.24(0.02)	1.20(0.03)
	1024	2.06(0.01)	1.37(0.02)	1.16(0.01)	1.06(0.02)
	2048	1.86(0.02)	1.02(0.01)	0.96(0.01)	1.06(0.01)
Blocks	512	2.53(9.13)	1.56(10.09)	0.16(2.91)	0.84(4.43)
	1024	2.30(10.93)	1.91(8.04)	0.14(1.58)	0.94(2.61)
	2048	1.63(3.94)	1.58(2.73)	0.10(0.98)	0.87(1.45)
Bumps	512	2.56(2.62)	2.62(2.71)	0.89(0.69)	0.73(0.48)
	1024	1.44(1.52)	2.07(1.55)	1.06(0.36)	0.91(0.31)
	2048	1.10(0.32)	1.02(0.30)	0.95(0.17)	0.85(0.20)
Corner	512	0.97(0.01)	1.07(0.01)	1.63(0.02)	1.16(0.02)
	1024	0.98(0.01)	1.05(0.01)	1.64(0.01)	1.10(0.01)
	2048	0.98(0.00)	1.03(0.01)	1.32(0.01)	1.06(0.01)
Doppler	512	1.14(0.27)	0.95(0.09)	1.11(0.10)	0.82(0.07)
	1024	1.06(0.10)	0.96(0.05)	1.24(0.05)	0.87(0.05)
	2048	0.87(0.04)	0.97(0.02)	1.32(0.03)	0.92(0.02)
Heavisine	512	0.98(3.37)	1.04(4.82)	1.60(5.60)	1.10(6.27)
	1024	1.01(2.76)	1.05(2.93)	1.94(3.35)	1.11(3.36)
	2048	0.98(1.45)	1.02(1.74)	1.50(2.12)	1.04(2.17)
Spikes	512	1.43(631.83)	0.96(85.78)	0.94(82.13)	0.97(98.93)
	1024	1.50(160.68)	0.96(45.71)	0.93(42.91)	0.93(49.32)
	2048	1.11(93.28)	0.98(23.47)	0.96(25.73)	0.97(29.33)
Wave	512	1.19(0.01)	0.96(0.02)	1.08(0.02)	0.67(0.02)
	1024	0.98(0.01)	1.03(0.01)	1.61(0.01)	0.98(0.01)
	2048	1.09(0.01)	1.02(0.00)	2.00(0.01)	1.07(0.01)

* Rows without a highlighted average MSE are cases in which VisuShrink had the lowest average MSE

Table 8: Haar basis with T_3 noise and SNR=5: Ratio of MSE of listed method to MSE of VisuShrink with corresponding standard deviation ($\times 10^{-3}$) of MSE of listed method*

Function	n	LD Block	Nason	UH	Ebayes
Blip	512	0.54(0.02)	0.45(0.06)	1.10(0.22)	1.36(0.23)
	1024	0.41(0.02)	0.43(0.03)	1.10(0.15)	1.43(0.16)
	2048	0.28(0.02)	0.31(0.03)	1.06(0.15)	1.46(0.15)
Blocks	512	1.71(19.32)	1.44(16.60)	0.72(25.04)	1.24(26.21)
	1024	1.01(2.99)	1.35(13.16)	0.82(17.11)	1.38(18.40)
	2048	0.62(2.42)	0.85(5.28)	0.86(16.71)	1.39(17.48)
Bumps	512	1.82(2.04)	2.87(3.59)	0.99(2.45)	1.09(2.51)
	1024	1.19(1.45)	2.23(3.71)	1.06(1.76)	1.21(1.83)
	2048	0.67(0.72)	1.41(2.57)	1.02(1.67)	1.23(1.72)
Corner	512	0.17(0.01)	0.29(0.01)	1.12(0.14)	1.31(0.16)
	1024	0.17(0.01)	0.22(0.01)	1.15(0.10)	1.42(0.11)
	2048	0.18(0.01)	0.19(0.01)	1.11(0.10)	1.45(0.10)
Doppler	512	0.64(0.24)	0.95(0.20)	1.05(0.49)	1.01(0.48)
	1024	0.30(0.11)	0.79(0.08)	1.11(0.32)	1.18(0.35)
	2048	0.34(0.03)	0.58(0.11)	1.12(0.31)	1.30(0.33)
Heavisine	512	0.17(4.28)	0.34(5.01)	1.16(50.53)	1.28(52.70)
	1024	0.20(2.82)	0.24(3.05)	1.20(34.44)	1.41(37.15)
	2048	0.19(2.77)	0.21(3.14)	1.13(33.48)	1.44(35.22)
Spikes	512	0.90(345.62)	1.03(754.27)	1.00(571.63)	1.18(596.81)
	1024	0.36(111.53)	0.83(348.20)	1.01(386.95)	1.29(411.52)
	2048	0.18(37.16)	0.58(100.40)	1.02(373.10)	1.37(397.86)
Wave	512	0.34(0.03)	0.77(0.03)	1.03(0.15)	1.02(0.15)
	1024	0.15(0.01)	0.41(0.01)	1.13(0.10)	1.22(0.10)
	2048	0.16(0.01)	0.25(0.01)	1.22(0.09)	1.43(0.10)

* Rows without a highlighted average MSE are cases in which VisuShrink had the lowest average MSE

To explore the distribution of λ by resolution level, we studied the distribution of 100 threshold values for each of four resolution levels when estimating the Blip function with T_3 noise and SNR=5. While the variance of λ was greater for coarse scales with fewer wavelet coefficients, the difference between the variance of λ at the highest resolution level and at the lowest resolution level was not significant. A visual comparison of the distributions is given in Figure 6 and shows comparable distributions of the threshold values across all resolution levels except $J - 3$. This can primarily be attributed to the choice of function, Blip, as evidenced in in Figure 7 where we observe the distributions of threshold values under the same conditions as before but for the Wave function and note comparable distributions across all levels. The similar pattern of outliers shown in Figures 6 and 7 indicates highly correlations between the values of λ for different resolution levels. As shown in Table 9, the Wave simulation described above yields high, positive correlations between threshold values, particularly for the three highest resolution levels.

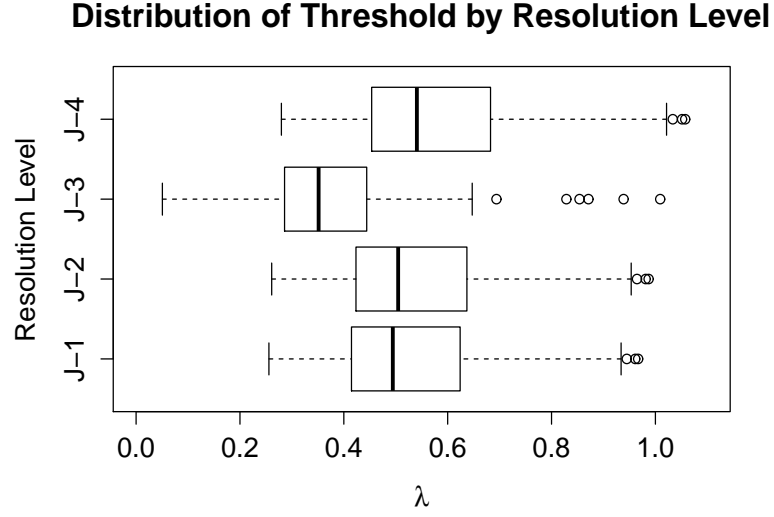


Figure 6: Distributions of λ by resolution level for 100 simulations of the Blip function with T_3 noise and SNR=5.




Figure 7: Distributions of λ by resolution level for 100 simulations of the Wave function with T_3 noise and SNR=5.

Table 9: Correlations between λ values from different resolution levels for 100 simulations of the Wave function with T_3 noise and SNR=5.

Resolution Level	$J - 1$	$J - 2$	$J - 3$	$J - 4$
$J - 1$	1.00			
$J - 2$	1.00	1.00		
$J - 3$	0.99	0.99	1.00	
$J - 4$	0.63	0.63	0.63	1.00

4.2. Examples

Respiratory Inductance Plethysmography (IP) is a method of evaluating pulmonary ventilation by measuring the variation in size of the chest and abdominal wall. If calibrated properly, the output voltage of the inductance plethysmograph is proportional to the change in volume of the body part under evaluation, in this case the lungs.

In Nason (1996), data from approximately 80 s of plethysmograph recordings made by the Department of Anaesthesia at the Bristol Royal Infirmary were used to compare various thresholding methods. Figure 8 shows the original IP data consisting of 4096 data points and Figure 9 shows the reconstructions of all 6 thresholding methods in our comparisons. In order to avoid the introduction of artificial singularities at the boundaries of our reconstruction (unlike our simulated data, our actual data is not periodic), we extended the original data at both the beginning and end by reflecting the data about the boundaries. For computation, the data was extended to the next dyadic, but only the reconstruction of the observed portion of the data is pictured. In the figures, the two main sets of regular oscillations correspond to normal breathing, while the disturbed behavior in the middle of the profile corresponds to the patient vomiting.




Figure 8: Inductance Plethysmography Data









Figure 9: Inductance Plethysmography Data Reconstruction

With the exception of UH, all reconstructions are visually similar. The Nason reconstruction appears to be the most appealing as LD Block does not threshold some high-resolution coefficients that may come as a result of noise at the beginning of the data. Our method retains 12.7% of the coefficients and sets the rest to zero while Nason’s method retains 6.3% of the coefficients. The threshold values for our LD Block method (with the lowest resolution level set to $J - 4$) were $\lambda_{J-1} = 0.20$, $\lambda_{J-2} = 0.46$, $\lambda_{J-3} = 0.01$, and $\lambda_{J-4} = 0.00$. The threshold value for Nason’s method was $\lambda = 0.44$. Since block thresholding thresholds the sum of squared coefficients in any block, in order to meaningfully compare the thresholds we must divide each LD Block threshold by the block size ($L=8$) then take the square root. This gives us $\lambda_{J-1}^* = 0.16$, $\lambda_{J-2}^* = 0.24$, $\lambda_{J-3}^* = 0.04$, and $\lambda_{J-4}^* = 0.00$. The reconstructions appear essentially the same, but since our threshold values are all smaller than Nason’s threshold value, we do end up with slightly more noise at the left extreme of the profile. We also note that with a threshold value of zero at the lowest resolution level, our method did not threshold any of the coefficients in that level.

In another relevant example with real data, we considered the vertical density profile data from Walker and Wright (2000). As a measure of quality control in the manufacturing of particle board, the density of a particle board is measured at fixed vertical depths between the two faces of the board. The density measurements for a given board comprise its *vertical density profile* (VDP). Figure 10 shows the original VDP data of a single board which consisted of 314 observations and Figure 11 shows the reconstructions of the VDP for each considered method. Similar to our handling of the non-periodic data in our first example, we addressed potential boundary issues by reflecting the data about both boundaries to extend original data to the next dyadic.




Figure 10: Vertical Density Profile Data

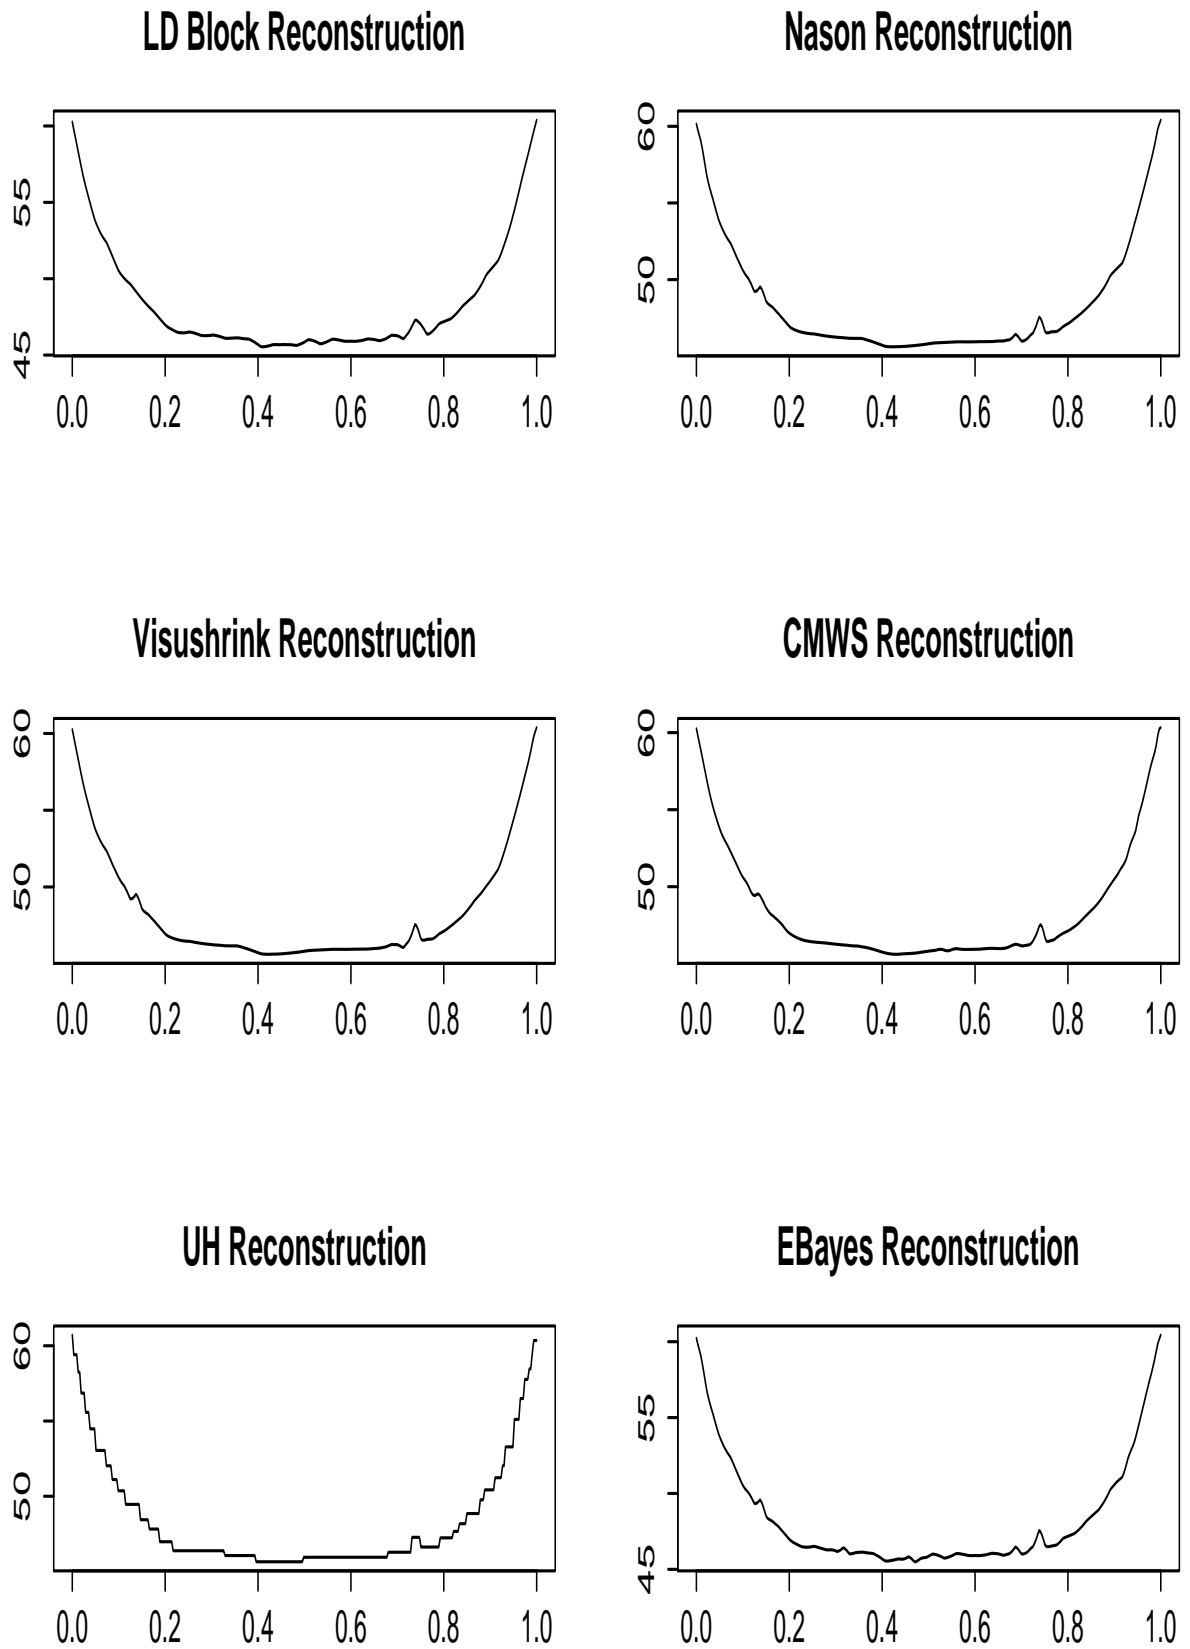


Figure 11: Vertical Density Profile Data Reconstruction

Again, with the exception of the more rigid reconstruction provided by the UH basis, reconstructions given in this example appear to be even more visually similar than in the last. There does appear to be a small peak around a vertical depth of 0.15 that our method thresholds and the other methods do not. Despite this observation, our method appears to retain a similar amount of detail throughout its reconstruction as it retains 12.5% of coefficients while Nason’s method retains only 10.6% of coefficients. The universal threshold value for Nason’s method was $\lambda = 0.46$. The thresholds from our method, after the same standardization for comparison described in our previous example, were $\lambda_{J-1}^* = 0.72$, $\lambda_{J-2}^* = 0.73$, $\lambda_{J-3}^* = 0.57$, and $\lambda_{J-4}^* = 0.25$.

5. Discussion and Conclusion

Our goal in developing the proposed method was to provide a completely non-parametric way to threshold wavelet coefficients when the distribution of the noise is unknown. Simulation results show that our level-dependent block thresholding method is an improvement over the existing cross-validated method as well as traditional methods for various types of non-Gaussian noise.

Other combinations of cross-validation and block thresholding methods were also considered. Generally speaking, global block thresholding (Block) was approximately equivalent to Nason’s method in terms of MSE, while the term-by-term level-dependent version of Nason’s method (LD) was an improvement over Nason’s global method (Nason) but not as good as the proposed level-dependent block method (LD Block). A version of the proposed method that was block-dependent as opposed to level-dependent (BD Block) was considered as well, but it did not perform very well. Since this method allows for a different threshold value to be chosen for every individual block, it was a very local analysis and was likely trying to model too much of the noise. It was also very computationally expensive. (See McGinnity and Chicken (2012) for full results). Thus the general MSE trend for non-normal noise was: LD Block < LD < BD Block.

An alternate cross-validation method was considered, in which we randomly chose (without replacement) half of the noisy data on which to perform the DWT and thresholding, and then compared that reconstruction to the other half of the data. However this method did not perform quite as well, and was significantly slower than the even-odd method, since numerous repetitions of the random selection and cross-validation were needed to ensure valid reconstructions.

When considering the type of noise to add on top of our test functions, we intentionally chose distributions with finite variance. This way the SNR is defined and can be used as a measure of the amount of noise in our data. Any distribution with infinite variance, such as Cauchy, cannot be properly scaled, and has arbitrarily large errors (see Figure 12).




Figure 12: The Doppler test function with Cauchy noise added to represent a “SNR” of 5.

The code and corresponding documentation to implement our method and generate the test functions used in this paper are provided as ancillary files for this article.

References

References

- Antoniadis, A., Fryzlewicz, P., 2006. Parametric modelling of thresholds across scales in wavelet regression. *Biometrika* 93, 465–471.
- Barber, S., Nason, G.P., 2004. Real nonparametric regression using complex wavelets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 66, 927–939.
- Cai, T., 1999. Adaptive wavelet estimation: A block thresholding and oracle inequality approach. *Ann. Statist.* 27, 898–924.
- Cai, T., Silverman, B., 2001. Incorporating information on neighboring coefficients into wavelet estimation. *Sankhya Ser. B* 63, 127–148.
- Chicken, E., 2003. Block thresholding and wavelet estimation for nonequispaced samples. *J. Statist. Plann. Inference* 116, 113–129.
- Chicken, E., 2005. Block-dependent thresholding in wavelet regression. *Journal of Nonparametric Statistics* 17, 467–491.
- Daubechies, I., 1992. *Ten Lectures on Wavelets*. SIAM, Philadelphia.
- Donoho, D., Johnstone, I., 1994. Ideal spatial adaptation via wavelet shrinkage. *Biometrika* 81, 425–455.
- Donoho, D., Johnstone, I., 1995. Adapting to unknown smoothness via wavelet shrinkage. *J. Am. Stat. Assoc.* 90, 1200–1224.
- Donoho, D., Johnstone, I., 1998. Minimax estimation via wavelet shrinkage. *Annals of Statistics* 26, 879–921.
- Downie, T.R., Silverman, B.W., 1996. The discrete multiple wavelet transform and thresholding methods. *IEEE Transactions in Signal Processing* 46, 2558–2561.
- Fryzlewicz, P., 2007. Unbalanced haar technique for nonparametric function estimation. *Journal of the American Statistical Association* 102, 1318–1327.
- Fryzlewicz, P., 2010. unbalhaar: Function estimation via Unbalanced Haar wavelets. URL: <https://cran.r-project.org/web/packages/unbalhaar/unbalhaar.pdf>. r package version 2.0.
- Girardi, M., Sweldens, W., 1997. A new class of unbalanced haar wavelets that form an unconditional basis for L_p on general measure spaces. *Journal of Fourier Analysis and Applications* 3, 457–474.
- Johnstone, I.M., Silverman, B.W., 2005. Empirical bayes selection of wavelet thresholds. *Annals of Statistics* , 1700–1752.
- Lawton, W., 1993. Applications of complex valued wavelet transforms to subband decomposition. *IEEE Transactions in Signal Processing* 41, 3566–3568.
- Lina, J.M., Mayrand, M., 1995. Complex daubechies wavelets. *Applied and Computaional Harmonic Analysis* 2, 219–229.
- Mallat, S., 1999. *A Wavelet Tour of Signal Processing*. 3 ed., Academic Press, San Diego.

- Mallat, S., Zhang, Z., 1993. Matching pursuit with time-frequency dictionaries. *IEEE Transactions in Signal Processing* 41, 3397–3415.
- Marron, S., Adak, S., Johnstone, I., Neuman, M., Patil, P., 1998. Exact risk analysis of wavelet regression. *J. Comput. Graph. Statist.* 7, 278–309.
- McGinnity, K., Chicken, E., 2012. Wavelet Block Thresholding for Non-Gaussian Errors. Technical Report M1006. Florida State University Department of Statistics. URL: <http://stat.fsu.edu/techreports/M1006.pdf>.
- Nason, G., 1996. Wavelet shrinkage using cross-validation. *Journal of the Royal Statistical Society* 58, 463–479.
- Nason, G., 1998. WaveThresh3 Software. Department of Mathematics, University of Bristol, Bristol, UK.
- Nason, G., 2013. wavethresh: Wavelets statistics and transforms. URL: <https://cran.r-project.org/web/packages/wavethresh/wavethresh.pdf>. r package version 4.6.6.
- Neumann, M.H., von Sachs, R., 1995. Wavelet thresholding: Beyond the Gaussian I.I.D. situation, in: *Lect. Notes Statist*, Springer-Verlag. pp. 301–329.
- Ogden, R., 1997. *Essential Wavelets for Statistical Applications and Data Analysis*. Birkhauser, Boston.
- Opsomer, J., Wang, Y., Yang, Y., 2001. Nonparametric regression with correlated errors. *Statistical Science* 16, 134–153.
- Pensky, M., 1999. Estimation of a smooth density function using meyer-type wavelets. *Statistics and Decisions* 17, 111–123.
- Pensky, M., Sapatinas, T., 2007. Frequentist optimatlity of Bayes factor estimators in wavelet regression models. *Statistica Sinica* 17, 599–633.
- R Development Core Team, 2009. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria. URL: <http://www.R-project.org>. ISBN 3-900051-07-0.
- Silverman, B.W., 2010. EbayesThresh: Empirical Bayes Thresholding and Related Methods. URL: <https://cran.r-project.org/web/packages/EbayesThresh/EbayesThresh.pdf>. r package version 1.3.2.
- Stein, C., 1981. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics* 9, 1135–1151.
- Vidakovic, B., 1999. *Statistical Modeling by Wavelets*. Wiley, New York.
- Walker, E., Wright, S.P., 2000. Comparing curves using additive models. *Journal of Quality Technology* 34, 118–129.
- Whitcher, B., 2009. waveslim: Basic wavelet routines for one-, two- and three-dimensional signal processing. URL: <http://CRAN.R-project.org/package=waveslim>. r package version 1.6.3.